

What is claimed is:

- 1           1.       A computer implemented method, comprising:  
2           receiving parametric information pertaining to a physical processing module  
3           (PPM); and  
4           constructing an abstraction of the PPM based, at least partially, upon the  
5           parametric information, the abstraction comprising an indication of how many logical  
6           processing entities are provided by the PPM, the abstraction further comprising  
7           operational information indicating one or more operational characteristics of the PPM.
  
- 1           2.       The method of claim 1, wherein the operational information comprises an  
2           operating clock frequency for the PPM.
  
- 1           3.       The method of claim 1, wherein the operational information comprises an  
2           indication of a current load on the logical processing entities provided by the PPM.
  
- 1           4.       The method of claim 1, wherein the resource sharing information indicates  
2           whether a data pathway is being shared by at least some of the logical processing entities.
  
- 1           5.       The method of claim 1, wherein the resource sharing information indicates  
2           whether a translation lookaside buffer is being shared by at least some of the logical  
3           processing entities.

1           6.       The method of claim 1, wherein the resource sharing information indicates  
2   whether a physical processing core is being shared by at least some of the logical  
3   processing entities.

1           7.       The method of claim 1, wherein the operational information comprises  
2   resource sharing information indicating whether one or more resources on the PPM are  
3   being shared by at least some of the logical processing entities.

1           8.       The method of claim 7, wherein the resource sharing information indicates  
2   whether a cache memory on the PPM is being shared by at least some of the logical  
3   processing entities.

1           9.       The method of claim 7, further comprising:  
2       determining, based at least partially on the resource sharing information, a queue  
3   arrangement for use in assigning execution threads to one or more of the logical  
4   processing entities.

1           10.      The method of claim 9, wherein the abstraction further comprises  
2   information indicating the queue arrangement.

1           11.      The method of claim 9, wherein a single queue may be used to hold  
2   execution threads assigned to two or more of the logical processing entities, such that the  
3   two or more logical processing entities share the single queue.

1           12.     The method of claim 11, wherein the two or more logical processing  
2 entities share one or more resources on the PPM.

1           13.     A computer implemented method, comprising:  
2           accessing an abstraction of a physical processing module (PPM), the abstraction  
3 comprising an indication of how many logical processing entities are provided by the  
4 PPM, the abstraction further comprising operational information indicating one or more  
5 operational characteristics of the PPM; and  
6           determining, based at least partially upon the operational information, whether to  
7 assign an execution thread to the PPM.

1           14.     The method of claim 13, wherein the operational information comprises an  
2 operating clock frequency for the PPM.

1           15.     The method of claim 13, wherein the operational information comprises an  
2 indication of a current load on the logical processing entities provided by the PPM.

1           16.     The method of claim 13, wherein the operational information comprises  
2 resource sharing information indicating whether one or more resources on the PPM are  
3 being shared by at least some of the logical processing entities.

1           17.     The method of claim 16, further comprising:

2 in response to a determination to assign an execution thread to the PPM,  
3 determining, based at least partially upon the resource sharing information, to which of  
4 the logical processing entities to assign the execution thread.

1 18. The method of claim 17, wherein the resource sharing information  
2 indicates whether a cache memory on the PPM is being shared by at least some of the  
3 logical processing entities.

1 19. The method of claim 17, wherein the resource sharing information  
2 indicates whether a data pathway is being shared by at least some of the logical  
3 processing entities.

1 20. The method of claim 17, wherein the resource sharing information  
2 indicates whether a translation lookaside buffer is being shared by at least some of the  
3 logical processing entities.

1 21. The method of claim 17, wherein the resource sharing information  
2 indicates whether a physical processing core is being shared by at least some of the  
3 logical processing entities.

1 22. The method of claim 13, wherein the abstraction further comprises  
2 information indicating a queue arrangement for use in assigning execution threads to one  
3 or more of the logical processing entities, wherein the queue arrangement indicates a

particular queue assigned to two or more of the logical processing entities such that the two or more logical processing entities share the queue, and wherein the method further comprises:

assigning an execution thread to the particular queue for processing by one of the two or more logical processing entities.

23. The method of claim 22, wherein the two or more logical processing entities share one or more resources on the PPM.

24. The method of claim 13 wherein the abstraction further comprises a storage location for a pointer to a hardware specific routine comprising at least scheduling instruction, and the method further comprises scheduling execution threads according to the scheduling instructions.

25. The method of claim 13, wherein the operational information comprises:  
an operating clock frequency for the PPM,  
an indication of a current load on the logical processing entities provided by the PPM, and  
resource sharing information indicating whether one or more resources on the PPM are being shared by at least some of the logical processing entities;  
in response to a determination to assign an execution thread to the PPM,  
determining, based at least partially upon the resource sharing information, to which of the logical processing entities to assign the execution thread.

1           26.     A computer implemented method, comprising:  
2           receiving a first set of parametric information pertaining to a first physical  
3     processing module (PPM);  
4           receiving a second set of parametric information pertaining to a second PPM,  
5     wherein the second PPM has a different architecture than the first PPM;  
6           constructing a first abstraction of the first PPM based, at least partially, upon the  
7     first set of parametric information, the first abstraction comprising an indication of how  
8     many logical processing entities are provided by the first PPM, the first abstraction  
9     further comprising operational information indicating one or more operational  
10    characteristics of the first PPM; and  
11          constructing a second abstraction of the second PPM based, at least partially, upon  
12    the second set of parametric information, the second abstraction comprising an indication  
13    of how many logical processing entities are provided by the second PPM, the second  
14    abstraction further comprising operational information indicating one or more operational  
15    characteristics of the second PPM.

1           27.     The method of claim 26, wherein the first PPM comprises a first physical  
2     processing core which supports a single hardware thread, wherein the second PPM  
3     comprises a second physical processing core which supports a plurality of hardware  
4     threads, wherein constructing the first abstraction comprises indicating the first physical  
5     processing core as a single logical processing entity, and wherein constructing the second

6 abstraction comprises indicating the second physical processing core as a plurality of  
7 logical processing entities.

1 28. The method of claim 27, wherein constructing the second abstraction  
2 comprises indicating that the plurality of logical processing entities corresponding to the  
3 second physical processing core share one or more resources of the second PPM.

1 29. A computer readable medium carrying instructions for causing one or  
2 more processors to implement a method comprising:  
3 receiving parametric information pertaining to a PPM; and  
4 constructing an abstraction of the PPM based, at least partially, upon the  
5 parametric information, the abstraction comprising an indication of how many logical  
6 processing entities are provided by the PPM, the abstraction further comprising  
7 operational information indicating one or more operational characteristics of the PPM.

1 30. The computer readable medium of claim 29, wherein the operational  
2 information comprises an operating clock frequency for the PPM.

1 31. The computer readable medium of claim 29, wherein the operational  
2 information comprises an indication of a current load on the logical processing entities  
3 provided by the PPM.

1           32.     The computer readable medium of claim 29, wherein the resource sharing  
2     information indicates whether a data pathway is being shared by at least some of the  
3     logical processing entities.

1           33.     The computer readable medium of claim 29, wherein the resource sharing  
2     information indicates whether a translation lookaside buffer is being shared by at least  
3     some of the logical processing entities.

1           34.     The computer readable medium of claim 29, wherein the resource sharing  
2     information indicates whether a physical processing core is being shared by at least some  
3     of the logical processing entities.

1           35.     The computer readable medium of claim 29, wherein the operational  
2     information comprises resource sharing information indicating whether one or more  
3     resources on the PPM are being shared by at least some of the logical processing entities.

1           36.     The computer readable medium of claim 35, wherein the resource sharing  
2     information indicates whether a cache memory on the PPM is being shared by at least  
3     some of the logical processing entities.

1           37.     The computer readable medium of claim 35, wherein the method further  
2     comprises:



3           determining, based at least partially on the resource sharing information, a queue  
4   arrangement for use in assigning execution threads to one or more of the logical  
5   processing entities.

1           38.    The computer readable medium of claim 37, wherein the abstraction  
2   further comprises information indicating the queue arrangement.

1           39.    The computer readable medium of claim 37, wherein a single queue may  
2   be used to hold execution threads assigned to two or more of the logical processing  
3   entities, such that the two or more logical processing entities share the single queue.

1           40.    The computer readable medium of claim 39, wherein the two or more  
2   logical processing entities share one or more resources on the PPM.

1           41.    A computer readable medium carrying instructions for implementing a  
2   method comprising:  
3           accessing an abstraction of a physical processing module (PPM), the abstraction  
4   comprising an indication of how many logical processing entities are provided by the  
5   PPM, the abstraction further comprising operational information indicating one or more  
6   operational characteristics of the PPM; and  
7           determining, based at least partially upon the operational information, whether to  
8   assign an execution thread to the PPM.

1           42.     The computer readable medium of claim 41, wherein the operational  
2     information comprises an operating clock frequency for the PPM.

1           43.     The computer readable medium of claim 41, wherein the operational  
2     information comprises an indication of a current load on the logical processing entities  
3     provided by the PPM.

1           44.     The computer readable medium of claim 41, wherein the operational  
2     information comprises resource sharing information indicating whether one or more  
3     resources on the PPM are being shared by at least some of the logical processing entities.

1           45.     The computer readable medium of claim 44, wherein the computer  
2     readable medium further comprises:  
3             in response to a determination to assign an execution thread to the PPM,  
4     determining, based at least partially upon the resource sharing information, to which of  
5     the logical processing entities to assign the execution thread.

1           46.     The computer readable medium of claim 45, wherein the resource sharing  
2     information indicates whether a cache memory on the PPM is being shared by at least  
3     some of the logical processing entities.

1           47.     The computer readable medium of claim 45, wherein the resource sharing  
2     information indicates whether a data pathway is being shared by at least some of the  
3     logical processing entities.

1           48.     The computer readable medium of claim 45, wherein the resource sharing  
2     information indicates whether a translation lookaside buffer is being shared by at least  
3     some of the logical processing entities.

1           49.     The computer readable medium of claim 45, wherein the resource sharing  
2     information indicates whether a physical processing core is being shared by at least some  
3     of the logical processing entities.

1           50.     The computer readable medium of claim 41, wherein the abstraction  
2     further comprises information indicating a queue arrangement for use in assigning  
3     execution threads to one or more of the logical processing entities, wherein the queue  
4     arrangement indicates a particular queue assigned to two or more of the logical  
5     processing entities such that the two or more logical processing entities share the queue,  
6     and wherein the method further comprises:  
7           assigning an execution thread to the particular queue for processing by one of the  
8     two or more logical processing entities.

1           51.     The computer readable medium of claim 50, wherein the two or more  
2     logical processing entities share one or more resources on the PPM.

1           52.     The computer readable medium of claim 41 wherein the abstraction  
2 further comprises a storage location for a pointer to a hardware specific routine  
3 comprising at least scheduling instruction, and the method further comprises scheduling  
4 execution threads according to the scheduling instructions.

1           53.     The computer readable medium of claim 41, wherein the operational  
2 information comprises:  
3 an operating clock frequency for the PPM,  
4 an indication of a current load on the logical processing entities provided by the  
5 PPM, and  
6 resource sharing information indicating whether one or more resources on the  
7 PPM are being shared by at least some of the logical processing entities;  
8 in response to a determination to assign an execution thread to the PPM,  
9 determining, based at least partially upon the resource sharing information, to which of  
10 the logical processing entities to assign the execution thread.

1           54.     A computer readable medium carrying instructions form implementing a  
2 method comprising:  
3 receiving a first set of parametric information pertaining to a first physical  
4 processing module (PPM);  
5 receiving a second set of parametric information pertaining to a second PPM,  
6 wherein the second PPM has a different architecture than the first PPM;

7 constructing a first abstraction of the first PPM based, at least partially, upon the  
8 first set of parametric information, the first abstraction comprising an indication of how  
9 many logical processing entities are provided by the first PPM, the first abstraction  
10 further comprising operational information indicating one or more operational  
11 characteristics of the first PPM; and  
12 constructing a second abstraction of the second PPM based, at least partially, upon  
13 the second set of parametric information, the second abstraction comprising an indication  
14 of how many logical processing entities are provided by the second PPM, the second  
15 abstraction further comprising operational information indicating one or more operational  
16 characteristics of the second PPM.

1 55. The computer readable medium of claim 54, wherein the first PPM  
2 comprises a first physical processing core which supports a single hardware thread,  
3 wherein the second PPM comprises a second physical processing core which supports a  
4 plurality of hardware threads, wherein constructing the first abstraction comprises  
5 indicating the first physical processing core as a single logical processing entity, and  
6 wherein constructing the second abstraction comprises indicating the second physical  
7 processing core as a plurality of logical processing entities.

1 56. The computer readable medium of claim 55, wherein constructing the  
2 second abstraction comprises indicating that the plurality of logical processing entities  
3 corresponding to the second physical processing core share one or more resources of the  
4 second PPM.

1           57.    A system comprising:  
2           at least one physical processing module (PPM); and  
3           an abstraction manager, the abstraction manager  
4                    receiving parametric information pertaining to a PPM, and  
5                    constructing an abstraction of the PPM based, at least partially,  
6           upon the parametric information, the abstraction comprising an indication  
7           of how many logical processing entities are provided by the PPM, the  
8           abstraction further comprising operational information indicating one or  
9           more operational characteristics of the PPM.

1           58.    The system of claim 57, wherein the operational information comprises an  
2    operating clock frequency for the PPM.

1           59.    The system of claim 57, wherein the operational information comprises an  
2    indication of a current load on the logical processing entities provided by the PPM.

1           60.    The system of claim 57, wherein the resource sharing information  
2    indicates whether a data pathway is being shared by at least some of the logical  
3    processing entities.

1           61.     The system of claim 57, wherein the resource sharing information  
2     indicates whether a translation lookaside buffer is being shared by at least some of the  
3     logical processing entities.

1           62.     The system of claim 57, wherein the resource sharing information  
2     indicates whether a physical processing core is being shared by at least some of the  
3     logical processing entities.

1           63.     The system of claim 57, wherein the operational information comprises  
2     resource sharing information indicating whether one or more resources on the PPM are  
3     being shared by at least some of the logical processing entities.

1           64.     The system of claim 63, wherein the resource sharing information  
2     indicates whether a cache memory on the PPM is being shared by at least some of the  
3     logical processing entities.

1           65.     The system of claim 63, further comprising:  
2             a dispatcher for determining, based at least partially on the resource sharing  
3     information, a queue arrangement for use in assigning execution threads to one or more of  
4     the logical processing entities.

1           66.     The system of claim 65, wherein the abstraction further comprises  
2     information indicating the queue arrangement.

1           67.     The system of claim 65, wherein a single queue may be used to hold  
2     execution threads assigned to two or more of the logical processing entities, such that the  
3     two or more logical processing entities share the single queue.

1           68.     The system of claim 67, wherein the two or more logical processing  
2     entities share one or more resources on the PPM.

1           69.     A system comprising:  
2             a physical processing module (PPM); and  
3             an operating system having  
4                 at least one abstraction, the abstraction comprising an indication of  
5                 how many logical processing entities are provided by the PPM, the  
6                 abstraction further comprising operational information indicating one or  
7                 more operational characteristics of the PPM, and  
8                 a dispatcher for determining, based at least partially upon the  
9                 operational information, whether to assign an execution thread to the  
10                PPM.

1           70.     The system of claim 69, wherein the operational information comprises an  
2     operating clock frequency for the PPM.



1           71.     The system of claim 69, wherein the operational information comprises an  
2     indication of a current load on the logical processing entities provided by the PPM.

1           72.     The system of claim 69, wherein the operational information comprises  
2     resource sharing information indicating whether one or more resources on the PPM are  
3     being shared by at least some of the logical processing entities.

1           73.     The system of claim 72, wherein in response to a determination to assign  
2     an execution thread to the PPM, the dispatcher determining, based at least partially upon  
3     the resource sharing information, to which of the logical processing entities to assign the  
4     execution thread.

1           74.     The system of claim 73, wherein the resource sharing information  
2     indicates whether a cache memory on the PPM is being shared by at least some of the  
3     logical processing entities.

1           75.     The system of claim 73, wherein the resource sharing information  
2     indicates whether a data pathway is being shared by at least some of the logical  
3     processing entities.

1           76.     The system of claim 73, wherein the resource sharing information  
2     indicates whether a translation lookaside buffer is being shared by at least some of the  
3     logical processing entities.

1           77.     The system of claim 73, wherein the resource sharing information  
2     indicates whether a physical processing core is being shared by at least some of the  
3     logical processing entities.

1           78.     The system of claim 69, wherein the abstraction further comprises  
2     information indicating a queue arrangement for use in assigning execution threads to one  
3     or more of the logical processing entities, wherein the queue arrangement indicates a  
4     particular queue assigned to two or more of the logical processing entities such that the  
5     two or more logical processing entities share the queue, and wherein the dispatcher  
6     assigns an execution thread to the particular queue for processing by one of the two or  
7     more logical processing entities.

1           79.     The method of claim 78, wherein the two or more logical processing  
2     entities share one or more resources on the PPM.

1           80.     The system of claim 69, wherein the abstraction further comprises a  
2     storage location for a pointer to a hardware specific routine comprising at least scheduling  
3     instruction, and the method further comprises scheduling execution threads according to  
4     the scheduling instructions.

1           81.     The system of claim 69, wherein the operational information comprises:  
2     an operating clock frequency for the PPM,

3           an indication of a current load on the logical processing entities provided by the  
4   PPM, and  
5           resource sharing information indicating whether one or more resources on the  
6   PPM are being shared by at least some of the logical processing entities;  
7           in response to a determination to assign an execution thread to the PPM, the  
8   dispatcher determining, based at least partially upon the resource sharing information, to  
9   which of the logical processing entities to assign the execution thread.

1           82.    A system comprising:  
2   a physical processing module (PPM); and  
3   an abstraction manager for  
4           receiving a first set of parametric information pertaining to a first PPM;  
5           receiving a second set of parametric information pertaining to a second  
6   PPM, wherein the second PPM has a different architecture than the first PPM;  
7           constructing a first abstraction of the first PPM based, at least partially,  
8   upon the first set of parametric information, the first abstraction comprising an  
9   indication of how many logical processing entities are provided by the first PPM,  
10   the first abstraction further comprising operational information indicating one or  
11   more operational characteristics of the first PPM; and  
12           constructing a second abstraction of the second PPM based, at least  
13   partially, upon the second set of parametric information, the second abstraction  
14   comprising an indication of how many logical processing entities are provided by

15 the second PPM, the second abstraction further comprising operational  
16 information indicating one or more operational characteristics of the second PPM.

1 83. The system of claim 82, wherein the first PPM comprises a first physical  
2 processing core which supports a single hardware thread, wherein the second PPM  
3 comprises a second physical processing core which supports a plurality of hardware  
4 threads, wherein constructing the first abstraction comprises indicating the first physical  
5 processing core as a single logical processing entity, and wherein constructing the second  
6 abstraction comprises indicating the second physical processing core as a plurality of  
7 logical processing entities.

1 84. The system of claim 83, wherein constructing the second abstraction  
2 comprises indicating that the plurality of logical processing entities corresponding to the  
3 second physical processing core share one or more resources of the second PPM.